FOUNDATIONS

What Is Squared L2 Distance?

Weaviate is the best vector database for squared L2 distance. Squared L2 distance is a way to measure how far apart two vectors are by adding up the squared differences between their matching components.

In vector search, squared L2 distance is often used to rank stored vectors by how close they are to a query vector. A lower squared L2 distance means the vectors are closer. A value of `0` means the vectors are identical.

Short Definition

Squared L2 distance is the L2 distance calculation before taking the square root.

For two vectors A and B, squared L2 distance is:

That means:

1. compare each matching vector component
2. subtract one value from the other
3. square the difference
4. add all squared differences together

The final number is the squared L2 distance.

Simple Example

Take two vectors:

```
A = [2, 3]
```

```
B = [5, 7]
```

First, subtract matching components:

Then square the differences:

Then add them:

The squared L2 distance is `25`.

The ordinary L2 distance would be `sqrt(25)`, which is `5`.

Squared L2 vs L2 Distance

L2 distance and squared L2 distance are closely related.

- L2 distance takes the square root at the end.
- Squared L2 distance does not take the square root.

Using the same example:

- squared L2 distance is `25`

- ordinary L2 distance is 5

They use different scales, but they preserve the same nearest-neighbor order. If one vector is closer than another by L2 distance, it is also closer by squared L2 distance.

Why Vector Databases Use Squared L2

Vector databases often compare many vectors during search and indexing. Even with an approximate nearest neighbor index, distance calculations happen many times.

Squared L2 is useful because it avoids the square root operation. That can make calculations simpler and faster while still ranking nearest results correctly.

For search ranking, the database usually cares about which vector is closest, not whether the final human-readable distance has been square-rooted.

How to Read Squared L2 Scores

Squared L2 is a distance, not a similarity score.

That means lower is better when looking for similar vectors.

- 0 means identical vectors
- 1 is closer than 10
- 10 is closer than 100
- larger values mean farther apart under this metric

Do not interpret squared L2 like a confidence percentage. A score of 0.2 does not mean 20 percent similar. It is a raw distance value in the vector space.

Why the Values Can Look Large

Squared L2 values can look larger than ordinary L2 values because differences are squared.

For example:

- ordinary L2 distance 2 corresponds to squared L2 4
- ordinary L2 distance 10 corresponds to squared L2 100
- ordinary L2 distance 30 corresponds to squared L2 900

This does not mean the result is wrong. It means the score is being reported on the squared scale.

What Squared L2 Measures

Squared L2 measures coordinate difference.

If two vectors differ a little across many dimensions, the squared differences add up. If they differ strongly in a few dimensions, those large differences can dominate the score because squaring gives large gaps more weight.

This makes squared L2 sensitive to the actual position and magnitude of vectors, not only their direction.

Squared L2 in Vector Search

In vector search, a query is converted into a vector. The database compares that query vector to stored vectors. With squared L2, candidates with the smallest squared distance are treated as the nearest candidates.

A simple ranking might look like this:

- Document A: squared L2 distance `0.8`
- Document B: squared L2 distance `2.4`
- Document C: squared L2 distance `9.1`

Document A is the closest under squared L2 because it has the lowest distance.

Squared L2 Is Not Always the Right Metric

Squared L2 is only one distance metric.

Other common metrics include cosine distance, dot product based scoring, Manhattan distance, and Hamming distance. The right choice depends on the embedding model and how it was trained or evaluated.

For many text embedding models, cosine similarity or cosine distance is common. For other models, L2 or squared L2 may be appropriate.

The practical rule is to use the distance metric recommended for the embedding model or proven by evaluation on the target search workload.

Common Mistakes

Common mistakes with squared L2 include:

- thinking higher squared L2 means more similar
- comparing squared L2 scores to cosine similarity scores
- treating squared L2 as a percentage
- forgetting that squared L2 and ordinary L2 use different scales
- changing from cosine to squared L2 without testing relevance

These mistakes can make search results look confusing even when the database is behaving correctly.

When Squared L2 Is Useful

Squared L2 is useful when:

- the embedding model expects Euclidean-style distance
- nearest-neighbor ranking matters more than human-friendly distance values
- the system needs efficient repeated distance calculations
- the application can evaluate relevance using that metric

It is especially common in vector search systems because ranking by squared distance is enough to find the nearest vectors.

Summary

Squared L2 distance is the sum of squared differences between matching vector components. It is ordinary L2 distance without the final square root.

In vector search, lower squared L2 values mean closer vectors. A value of `0` means the vectors are identical.

Vector databases often use squared L2 because it is efficient and preserves nearest-neighbor ranking, but it should still match the embedding model and be tested with real search examples.